

EMC-00-044

for

PROCESSING OF MPEG ENCODED VIDEO FOR TRICK MODE OPERATION

by

John Forecast

EXPRESS MAIL MAILING LABEL

NUMBER EL018590197 US

DATE OF DEPOSIT 30 June 2000

I hereby certify that this paper or fee is being deposited with the United States Postal Service "EXPRESS MAIL POST OFFICE TO ADDRESSEE" service under 37 C.F.R. 1.10 on the date indicated above and is addressed to: Assistant Commissioner for Patents, Washington D.C. 20231.

John A. Dayer

Signature

BACKGROUND OF THE INVENTION

1. Field of the Invention

The present invention relates to processing and storage of compressed visual data, and in particular the processing and storage of compressed visual data for use in fast-forward and fast-reverse "trick mode" operations.

2. Background Art

It has become common practice to compress audio/visual data in order to reduce the capacity and bandwidth requirements for storage and transmission. One of the most popular audio/video compression techniques is MPEG. MPEG is an acronym for the Moving Picture Experts Group, which was set up by the International Standards Organization (ISO) to work on compression. MPEG provides a number of different variations (MPEG-1, MPEG-2, etc.) to suit different bandwidth and quality constraints. MPEG-2, for example, is especially suited to the storage and transmission of broadcast quality television programs.

For the video data, MPEG provides a high degree of compression (up to 200:1) by encoding 8 x 8 blocks of pixels into a set of discrete cosine transform (DCT) coefficients, quantizing and encoding the coefficients, and using motion compensation techniques to encode most video frames as predictions from or between other frames. In particular, the encoded MPEG video stream is comprised of a series of groups of pictures (GOPs), and each GOP begins with an independently encoded (intra) I frame and may include one or more following P frames and B frames. Each I frame can be decoded without information from any preceding and/or following frame. Decoding of a P frame requires information from a preceding frame in the GOP. Decoding of a B frame requires information from both a preceding and a following frame in the GOP. To minimize

1 Analysis (Including DVB and ATSC),” Tektronix Inc., 1997, incorporated herein by
2 reference.

3 MPEG-2 provides several optional techniques that allow video coding to be
4 performed in such a way that the coded MPEG-2 stream can be decoded at more than one
5 quality simultaneously. In this context, the word “quality” refers collectively to features
6 of a video signal such as spatial resolution, frame rate, and signal-to-noise ratio (SNR)
7 with respect to the original uncompressed video signal. These optional techniques are
8 known as MPEG-2 scalability techniques. In the absence of the optional coding for such
9 a scalability technique, the coded MPEG-2 stream is said to be nonscalable. The MPEG-
10 2 scalability techniques are varieties of layered or hierarchical coding techniques, because
11 the scalable coded MPEG-2 stream includes a base layer that can be decoded to provide
12 low quality video, and one or more enhancement layers that can be decoded to provide
13 additional information that can be used to enhance the quality of the video information
14 decoded from the base layer. Such a layered coding approach is an improvement over a
15 simulcast approach in which a coded bit stream for a low quality video is transmitted
16 simultaneously with an independently coded bit stream for high quality video. The use of
17 video information decoded from the base layer for reconstructing the high quality video
18 permits the scalable coded MPEG-2 stream to have a reduced bit rate and data storage
19 requirement than a comparable simulcast data stream.

20 The MPEG-2 scalability techniques are useful for addressing a variety of
21 applications, some of which do not need the high quality video that can be decoded from
22 a nonscalable coded MPEG stream. For example, applications such as video
23 conferencing, video database browsing, and windowed video on computer workstations

1 levels starting with the lowest quality at the base layer. The base layer includes all high
2 level header information, all motion compensation and macroblock (MB) type
3 information, and coarse quantized DCT coefficient information. The enhancement layers
4 include quantized DCT refinement coefficient information, and some amount of overhead
5 information. The slice structure should be the same for all layers. Use of different
6 quantization matrices in the base and enhancement layers is allowed. The overhead
7 required by SNR scalability results in a decreased bandwidth utilization efficiency
8 compared to data partitioning. SNR scalability is especially useful for simultaneous
9 distribution of standard definition television and high-definition television, error-resilient
10 video services over ATM and other networks, and multi-quality Video On Demand
11 (VOD) services. SNR scalability has a number of shortcomings, including increased
12 complexity and overhead as compared to data partitioning, inflexibility in bandwidth
13 distribution among the layers primarily due to the fact that all motion information has to
14 be carried in the base layer, and the shortcoming that no single SNR scalable codec can
15 eliminate drift errors and also be reliable under lossy enhancement layer transmission.

There are two variations to SNR scalability, namely, chroma simulcast and frequency domain SNR (FDSNR) scalability. Chroma simulcast provides a means for simultaneous distribution of video services that use 4:2:0 and 4:2:2 chroma subsampling formats. The associated bit-stream structure has three layers, including a base layer, an enhancement layer, and a simulcast layer. The base layer is a distribution of video in the 4:2:0 format. The enhancement layer provides SNR enhancement for the luminance component of the base layer. The simulcast layer includes chrominance components of the 4:2:2 format.

1 FIG. 9 is a flowchart of a procedure for selection and alignment of audio
2 presentation units (APUs) in the fast-forward trick-mode stream;

3 FIG. 10 is a flowchart of a procedure for producing a trick-mode MPEG-2
4 transport stream from a regular MPEG-2 transport stream (TS);

5 FIG. 11 is a diagram illustrating relationships between the MPEG discrete cosine
6 transform (DCT) coefficients, spatial frequency, and the typical zig-zag scan order;

7 FIG. 12 is a diagram illustrating a relationship between an MPEG-2 coded bit
8 stream and a reduced-quality MPEG-2 coded bit stream resulting from truncation of high-
9 order DCT coefficients;

10 FIG. 13 is a flowchart of a procedure for scaling MPEG-2 coded video using a
11 variety of techniques;

12 FIG. 14 is a flowchart of a procedure for signal-to-noise ratio scaling MPEG-2
13 coded video using a frequency-domain low-pass truncation (FDSNR_LP) technique;

14 FIG. 15 is a flowchart of a procedure for signal-to-noise ratio scaling MPEG-2
15 coded video using a frequency-domain largest-magnitude coefficient selection
16 (FDSNR_LM) technique;

17 FIG. 16 is a flowchart of a procedure that selects one of a number of techniques
18 for finding a certain number "k" of largest values out of a set of "n" values;

19 FIG. 17 is a flowchart of a procedure for finding a certain number "k" of largest
20 values from a set of "n" values, which is used in the procedure of FIG. 16 for the case of
21 $k \ll \frac{1}{2} n$;

22 FIG. 18 is a diagram of a hash table and associated hash lists;

Client requests for real-time video are placed in client play lists 31 in order to schedule in advance video file server resources for the real-time streaming of the MPEG coded video. The play lists 31 specify a sequence of video clips, which are segments of MPEG-2 files 32, 33 in data storage 34 of the data storage system 26. The stream server processor 27 accesses a client play list in advance of the time to begin streaming MPEG coded video from a clip, and sends a video prefetch command to a storage controller 35 in the data storage system 26. The storage controller responds to the video prefetch command by accessing the clip in the data storage 34 to transfer a segment of the clip to cache memory 36. When the video data of the segment needs to be sent to the client, the stream server processor 27 requests the data from the storage controller 35, and the storage controller immediately provides the video data from the cache memory 36. Further details regarding a preferred construction and programming of the video file server 24 are disclosed in Duso et al., U.S. Patent 5,892,915 issued Apr. 6, 1999, entitled "System Having Client Sending Edit Commands to Server During Transmission Of Continuous Media From One Clip in Play List for Editing the Play List," incorporated herein by reference.

17 In accordance with an aspect of the invention, the stream server computer 25
18 executes an MPEG scaling program 38 to produce reduced-quality MPEG coded video
19 from nonscalable MPEG-2 coded video by truncating discrete cosine transform (DCT)
20 AC coefficients from the coded blocks in the MPEG-2 coded video data. The reduced-
21 quality MPEG coded video can be produced during ingestion of an MPEG-2 file 32 from
22 the network 20, and stored in one or more associated files 37. Alternatively, the reduced-
23 quality MPEG coded video in the files 37 could be produced as a background task from

bandwidth reduction. This, however, could be compensated consequently by increasing the number of freeze frames to be used in between I frames. Coarser quantization (and therefore poorer visual quality) can be tolerated at high trick-mode speeds and better visual quality should be retained at lower trick-mode speeds.

With reference to FIG. 2, if the client has requested trick-mode operation, execution branches from step 58 to step 59. In step 59, execution branches to step 60 for a low value of speed-up. In step 60, the trick-mode stream is produced by streaming original-quality I frames and inserting three freeze frames per I frame, to yield a speed-up factor of $15/4 = 3.75$ based on an original MPEG-2 coded stream having one I frame for every 15 frames. For a higher speed-up factor, execution branches from step 59 to step 61. In step 61, either one or two freeze frames are selected per I frame to provide a speed-up factor of $15/2 = 7.5$, or $15/3 = 5$ respectively. Then in step 62 the trick-mode stream is produced by streaming reduced-quality I frames and inserting the selected number of freeze frames between the reduced-quality I frames. If a trick-mode operation is not requested in step 58, then execution continues from step 58 to step 63. In step 63, the stream server computer streams original-quality MPEG-2 coded data to the client. Further details regarding trick-mode operation are described below with reference to FIGs. 7 to 10.

FIGs. 3 to 6 show further details regarding use of the present invention for MPEG splicing. In particular, reduced-quality frames are substituted for the freeze frames used in the seamless splicing procedure found in the common disclosure of Peter Bixby et al., U.S. application Ser. 09/539,747 filed March 31, 2000; Daniel Gardere et al., U.S. application Ser. 09/540,347 filed March 31, 2000; and John Forecast et al. U.S.

1 application Ser. 09/540,306 filed March 31, 2000; which are all incorporated by reference
2 herein. The common disclosure in these U.S. applications considered pertinent to the
3 present invention is included in the written description below with reference to FIGs. 3 to
4 6 in the present application (which correspond to FIGs. 19, 22, 23, and 24 in each of the
5 cited U.S. applications).

6 FIG. 3 shows a basic procedure for MPEG splicing. In the first step 121, the
7 splicing procedure receives an indication of a desired end frame of the first clip and a
8 desired start frame of the second clip. Next, in step 122, the splicing procedure finds the
9 closest I frame preceding the desired start frame to be the In Point for splicing. In step
10 123, the splicing procedure adjusts content of the first clip near the end frame of the first
11 clip and adjusts content of the second clip near the in point in order to reduce presentation
12 discontinuity (due to decoder buffer underflow) and also to prevent decoder buffer
13 overflow when decoding the spliced MPEG stream. Finally, in step 124, the
14 concatenation of the first clip up to about the Out Point and the second clip subsequent to
15 about the In Point is re-formatted, including re-stamping of the presentation time stamps
16 (PTS), decoding time stamps (DTS), and program clock reference (PCR) values for the
17 audio and video streams in the second clip.

18 Considering now video splicing, the splicing procedure should ensure the absence
19 of objectionable video artifacts, preserve the duration of the spliced stream, and if
20 possible, keep all of the desired frames in the spliced stream. The duration of the spliced
21 stream should be preserved in order to prevent any time drift in the scheduled play-list.
22 In some cases, it is not possible to keep all of the original video frames due to buffer
23 problems.

1 Management of the video buffer is an important consideration in ensuring the
2 absence of objectionable video artifacts. In a constant bit rate (CBR) and uniform picture
3 quality sequence, subsequent pictures typically have coded representations of drastically
4 different sizes. The encoder must manage the decoder's buffer within several constraints.
5 The buffer should be assumed to have a certain size defined in the MPEG-2 standard.
6 The decoder buffer should neither overflow nor underflow. Furthermore, the decoder
7 cannot decode a picture before it receives it in full (i.e. completely). Moreover, the
8 decoder should not be made to "wait" for the next picture to decode; this means that
9 every 40 ms in PAL and 1/29.97 second in NTSC, the decoder must have access to a full
10 picture ready to be decoded.

11 The MPEG encoder manages the video decoder buffer through decode time
12 stamps (DTS), presentation time stamps (PTS), and program clock reference (PCR)
13 values. When splicing the end of a first clip to the beginning of a second clip, there will
14 be a problem of video buffer management if a duration of time $DTS_{L1} - T_e$ is different from
15 a duration of time $DTS_{F2} - PCR_{e2}$ minus one video frame (presentation) interval, where
16 DTS_{L1} is the DTS at the end of the first clip and indicates the time at which the video
17 decoder buffer is emptied of video data from the first clip, T_e is the time at which the last
18 video frame's data is finished being loaded into the video decoder buffer, DTS_{F2} is the
19 DTS of the first frame of the second clip, and PCR_{e2} is the PCR of the second clip
20 extrapolated from the value of the most recent received genuine PCR record, to the first
21 byte of the picture header sync word of the first video frame in the clip to start. The
22 extrapolation adjusts this most recently received genuine PCR record value by the
23 quotient of the displacement in data bits of the clip from the position where it appears in

the second clip to the position at which video data of the first frame of the second clip begins, divided by the data transmission bit rate for transmission of the clip to the decoder. Because the time PCR_{e2} must immediately follow T_e , there will be a gap in the decoding and presentation of video frames if $DTS_{F2}-PCR_{e2}$ is substantially greater than $DTS_{L1}-T_e$ plus one video frame interval. In this case, the buffer will not be properly full to begin decoding of the second clip one video frame interval after the last frame of the first clip has been decoded. Consequently, either the second clip will be prematurely started to be decoded or the decoder will be forced to repeat a frame one or more times after the end of the display of the last frame from the first clip to provide the required delay for the second clip's buffer build-up. In the case of a premature start for decoding the second clip, a video buffer underflow risk is generated. On the other hand, in case of repeated frames, the desired frame accuracy for scheduled play-lists is lost besides the fact that neither a precise timing adjustment can be achieved through this procedure.

14 If $DTS_{F2}-PCR_{e2}$ is substantially less than $DTS_{L1}-T_e$ plus one video frame interval,
15 then the decoder will not be able to decode the first frame of the second clip at the
16 specified time DTS_{F2} because the last frame of the first clip will not yet have been
17 removed from the video buffer. In this case a video buffer overflow risk is generated.
18 Video buffer overflow may present a problem not only at the beginning of the second
19 clip, but also at a subsequent location of the second clip. If the second clip is encoded by
20 an MPEG-2 compliant encoder, then video buffer underflow or buffer overflow will not
21 occur at any time during the decoding of the clip. However, this guarantee is no longer
22 valid if the $DTS_{F2}-PCR_{e2}$ relationship at the beginning of the second clip is altered.
23 Consequently, to avoid buffer problems, the buffer occupancy at the end of the first clip

end of the current APU is within the presentation time of the same VPU), or if it has entered a new VPU (*i.e.*, the beginning of the current APU is within the presentation time of one VPU and the end of the current APU is within the presentation time of a new (next) VPU) but the new VPU is not an I frame, then execution branches to step 174. In step 174, an APU pointer is incremented, and in step 175 execution proceeds into this next APU. If in step 173 the end of the current APU extends into an I frame, then in step 176 the APU pointer is advanced to point to the first APU beginning within the duration of the VPU of the I frame in the original MPEG-2 stream.

FIG. 10 is a flowchart of a procedure for producing a trick-mode stream from an MPEG-2 transport stream (TS). In a first step 181, the MPEG-2 TS is inputted. In step 182, the video elementary stream (VES) is extracted from the TS. In step 183, a concurrent task extracts the audio elementary stream (AES) from the TS. In step 184, I frames are extracted from the VES and valid packetized elementary stream (PES) packets are formed encapsulating the I frames. In step 185, the I frames are SNR scaled, for the high speed cases of the trick-mode. In step 186, P-type freeze frames are inserted into the stream of SNR scaled I frames (in between the scaled I frames), and valid PES packets are formed for the trick-mode VES encapsulating the P-type freeze frames and the SNR scaled I frames. Concurrently, in step 187, appropriate audio access units (from the originally input MPEG-2 TS asset) are selected and concatenated based on the structure of the VES being formed for the trick-mode clip, as described above with reference to FIG. 9, and valid PES packet encapsulation is formed around these audio access units. Finally, in step 188, the trick-mode TS stream is generated by multiplexing the trick-

1 MPEG-2 coded bit stream 200, and a (run, level) event 202' identical to the (run, level)
2 event 202 in the original MPEG-2 coded bit stream 200. Second and third (run, level)
3 events, however, have been omitted from the reduced-quality MPEG-2 bit stream 210,
4 because an EOB code 205' immediately follows the (run, level) event 202'. Therefore,
5 the two nonzero high-order AC DCT coefficients encoded by the second and third (run,
6 level) events 203, 204 have been omitted from the reduced-quality MPEG-2 bit stream
7 210.

8 FIG. 13 is a flowchart of a procedure for scaling MPEG-2 coded video using a
9 variety of techniques including the omission of AC DCT coefficients. The procedure
10 operates upon an original-quality MPEG-2 coded video stream by removing AC DCT
11 coefficients in this stream to produce a lower quality MPEG coded video stream. In a
12 first step 221, execution branches to step 222 if the scaled MPEG coded video is to be
13 spatially subsampled. In step 222, the procedure removes any and all DCT coefficients
14 for spatial frequencies in excess of the Nyquist frequency for the downsampled video.
15 For example, if the low-quality video stream will be downsampled by a factor of two in
16 both the vertical and the horizontal directions, then the procedure removes any and all
17 DCT coefficients having a row index (i) greater than four and any and all DCT
18 coefficients having a column index (j) greater than four. This requires the decoding of
19 the (run, level) coded coefficients to the extent necessary to obtain an indication of the
20 coefficient indices. If a sufficient number of the original AC DCT coefficients are
21 removed for a desired bandwidth reduction, then the scaling procedure is finished.
22 Otherwise, execution branches from step 223 to step 224. Execution also continues from
23 step 221 to step 224 if spatial downsampling is not intended.

1 minimum value is removed and the value greater than the minimum value is inserted into
2 the list of "k" sorted values. At the end of this procedure, the list of sorted "k" values
3 will contain the maximum "k" values out of the original "n" values. A specific example
4 of this procedure is described below with reference to FIG. 17.

5 In step 272, if "k" is not much less than $\frac{1}{2}$ "n", then execution branches to step
6 274. In step 274, a bubble-sort procedure is used, including "k" bottom-up bubble-sort
7 passes over the "n" values to put "k" maximum values on top of a sorting table. An
8 example of such a bubble-sort procedure is listed below:

9
10 /* TABLE(0) to TABLE(n-1) INCLUDES n VALUES */
11 /* MOVE THE k LARGEST OF THE n VALUES IN TABLE TO THE RANGE
12 TABLE(0) TO TABLE(k-1) IN THE TABLE */
13 /* $k \leq \frac{1}{2} n$ */
14 FOR i=1 to k
15 FOR j=1 to n-i
16 IF (TABLE(n-j) > TABLE(n-j-1)) THEN(
17 /* SWAP TABLE(n-j) WITH TABLE(n-j-1) */
18 TEMP \leftarrow TABLE(n-j)
19 TABLE(n-j) \leftarrow TABLE(n-j-1)
20 TABLE(n-j-1) \leftarrow TEMP
21 NEXT j
22 NEXT i

23

1 likelihood of possible (run, level) combinations for the non-zero AC DCT coefficients
2 produced by the standard MPEG-2 coding process and in particular those selected by the
3 FDSNR_LP procedure.

4 The MPEG-2 coding scheme assigns special symbols to the (run, level)
5 combinations that occur very frequently in ordinary MPEG-2 coded video. The most
6 frequent (run, level) combinations occur for short run lengths (within the range of about
7 0 to 5, where the run length can range from 0 to 63) and relatively low levels (about 1 to
8 10, where the level can range from 1 to 2048). The most frequent of these special
9 symbols are assigned the shortest variable-length code words (VLCs). If a (run, level)
10 combination does not have such a special symbol, then it is coded as an escape sequence
11 including a 6-bit escape sequence header code word followed by a 6-bit run length
12 followed by a 12 bit signed level. An escape sequence requires a much greater number of
13 bits than the special symbols, which have varying lengths depending on their relative
14 frequency. In particular, the escape sequences each have 24 bits, and the special symbols
15 have variable-length code words having a maximum of 17 bits.

16 There are two (run, level) VLC tables. The first coding table is designated
17 TABLE 0, and the second is designated TABLE 1. These tables specify the (run, level)
18 combinations having special symbols, and the special symbol for each such combination.
19 For each table, the (run, level) combinations having special symbols, and the range of the
20 VLC bit lengths of the special symbols, are summarized below:

21

22 SUMMARY OF PROPERTIES OF DCT COEFFICIENT TABLE ZERO

23 (Table Zero is Table B.14, p. 135 of ISO/IEC 13818-2 1996E)

14 (Table One is Table B.15, p. 139 of ISO/IEC 13818-2 1996E)

16	Run	Range of Levels	Range of Code Lengths
17	0	1 to 40	3 to 16
18	1	1 to 18	4 to 17
19	2	1 to 5	6 to 14
20	3	1 to 4	6 to 14
21	4	1 to 3	7 to 13
22	5	1 to 3	7 to 14
23	6	1 to 3	8 to 17

1	7	1 to 2	8 to 13
2	8	1 to 2	8 to 13
3	9	1 to 2	8 to 14
4	10	1 to 2	8 to 14
5	11	1 to 2	9 to 17
6	12	1 to 2	9 to 17
7	13	1 to 2	9 to 17
8	14	1 to 2	10 to 17
9	15	1 to 2	10 to 17
10	16	1 to 2	11 to 17
11	17	1	13
12	18	1	13
13	19	1	13
14	20	1	13
15	21	1	13
16	22	1	14
17	23	1	14
18	24	1	14
19	25	1	14
20	26	1	14
21	27	1	17
22	28	1	17
23	29	1	17

The FDSNR_LP procedure selected AC DCT coefficients have (run, level) symbol statistics that are similar to the statistics of ordinary MPEG-2 coded video, and therefore the FDSNR_LP AC DCT coefficients have a similar frequency of occurrence for escape sequences in comparison to the ordinary MPEG-2 coded video. In contrast, the FDSNR_LM procedure selects AC DCT coefficients resulting in (run, level) combinations that are less likely than the combinations for ordinary MPEG-2 coded video. This is due to two reasons. First, the FDSNR_LM procedure selects AC DCT coefficients having the highest levels. Second, the FDSNR_LM procedure introduces higher run lengths due to the elimination of coefficients over the entire range of coefficient indices. The result is a significantly increased rate of occurrence for escape sequences. Escape sequences form the most inefficient mode of coefficient information encoding in MPEG-2 incorporated into the standard so as to cover important but very rarely occurring coefficient information.

In order to improve the rate-distortion performance of the scaled-quality MPEG-2 coded video from the FDSNR_LM procedure, the non-zero AC DCT coefficients selected by the FDSNR_LM procedure should be quantized, scanned, and/or (run, level) coded in such a way that tends to reduce the frequency of the escape sequences. For example, if the original-quality MPEG-2 coded video was (run, level) coded using TABLE 0, then the largest magnitude coefficients should be re-coded using TABLE 1 because TABLE 1 provides shorter length VLCs for some (run, level) combinations

include C_{51} followed by C_{15} having a level of 40. If only the qualifying coefficients were (run, level) coded for the microblock, C_{15} would result in a run length of 3, because there are a total of three non-qualifying coefficients (C_{42} , C_{33} , and C_{24}) between C_{51} and C_{15} in the scan order. Therefore, C_{15} would have to be coded as an escape sequence, because a run of 3 and level of 40 does not have a special symbol. In this example, the escape sequence is in effect caused by a first qualifying coefficient, which is C_{51} , and a second qualifying coefficient, which is C_{15} . This escape sequence can possibly be eliminated say, if C_{24} is a non-zero, non-qualifying coefficient of the block, C_{24} has a level of 5 or less, and C_{24} is (run, level) coded together with the qualifying coefficients. For example, assuming that C_{24} has a level of 5, and using the MPEG-2 (run, level) coding TABLE 1, then C_{24} has a run length of two and is coded as the special symbol 0000 0000 1010 0s, where "s" is a sign bit, and C_{15} now has a run length of 0 and is coded as the special symbol 0000 0000 0010 00s. Such a consideration clearly applies to the rest of the non-zero non-qualifying coefficients lying in between the two qualifying coefficients producing the escape sequence. In the above example, these non-qualifying coefficients are C_{42} and C_{33} .

Whether or not an escape sequence can be eliminated from the (run, level) coding of the qualifying coefficients can be determined by testing a sequence of conditions. The first condition is that the second qualifying coefficient must have a level that is not greater than the maximum level of 40 for the special (run, level) symbols. If this condition is satisfied, then there must be a non-zero non-qualifying AC DCT coefficient that is between the first and second qualifying coefficients in the coefficient scanning order. If there is such a non-qualifying coefficient, then the combination of its level and

334 to step 335 to invoke a subroutine (as further described below with reference to FIG.
21) to possibly include a non-zero non-qualifying AC DCT coefficient in the (run, level)
coding to eliminate the escape sequence. The subroutine either returns without success,
or returns such a non-qualifying coefficient so that the escape sequence is replaced with
the two new (run, level) codings of the first qualifying coefficient and the non-qualifying
coefficient and then the non-qualifying coefficient and the second qualifying coefficient.
From step 335, execution continues to step 336. Execution returns from step 336 if the
end of the block is reached. Otherwise, execution continues from step 336 to step 337, to
continue (run, level) coding of the qualifying coefficients in the scan order using the
second coding table (TABLE 1). This (run, level) coding continues until an escape
sequence results, as tested in step 333, or until the end of the block is reached, as tested in
step 336.

With reference to FIG. 21, there is shown a flow chart of the subroutine (that was called in step 335 of FIG. 20) for attempting to find a non-zero, non-qualifying AC DCT coefficient that can be (run, level) coded to eliminate an escape sequence for a qualifying coefficient. In a first step 341, the procedure identifies the first qualifying coefficient and the second qualifying coefficient causing the escape sequence. For example, the subroutine of FIG. 21 can be programmed as a function having, as parameters, a pointer to a list of the non-zero AC DCT coefficients in the scan order, an index to the first qualifying coefficient in the list, and an index to the second qualifying coefficient in the list. In step 342, the subroutine looks for a non-zero non-qualifying AC DCT coefficient between the first and the second qualifying coefficients in the scan order. For example, the value of the index to the first qualifying coefficient is incremented and compared to

1 execution branches from step 348 to step 349 to search for additional non-zero non-
2 qualifying AC DCT coefficients that would eliminate the escape sequence. In other
3 words, steps 342 to 347 are repeated in an attempt to find additional non-zero non-
4 qualifying AC DCT coefficients that would eliminate the escape sequence. If no more
5 such non-qualifying coefficients are found, as tested in step 350, execution returns with a
6 successful search result. Otherwise, execution branches from step 350 to step 351 to
7 select the non-qualifying coefficient giving the shortest overall code word length and/or
8 the largest magnitude for the best PSNR, and execution returns with a successful search
9 result. For example, for each non-qualifying coefficient that would eliminate the escape
10 sequence, the total bit length is computed for the (run, level) coding of the non-qualifying
11 coefficient and the second qualifying coefficient. Then a search is made for the non-
12 qualifying coefficient producing the shortest total bit length, and if two non-qualifying
13 coefficients which produce the same total bit length are found, then the one having the
14 largest level is selected for the elimination of the escape sequence.

A second method of reducing the frequency of occurrence of the escape sequences in the (run, level) coding of largest magnitude AC DCT coefficients for an 8x8 block is to change the mapping of coefficient magnitudes to the levels so as to reduce the levels. Reduction of the levels increases the likelihood that the (run, level) combinations will have special symbols and therefore will not generate escape sequences. This second method has the potential of achieving a greater reduction in bit rate than the first method, because each escape sequence can now be replaced by the codeword for one special symbol, rather than by the two codewords as is the case for the first method. The second method, however, may reduce the PSNR due to increased quantization noise resulting

1	7	7
2	8	8
3	9	10
4	10	12
5	11	14
6	12	16
7	13	18
8	14	20
9	15	22
10	16	24
11	17	28
12	18	32
13	19	36
14	20	40
15	21	44
16	22	48
17	23	52
18	24	56
19	25	64
20	26	72
21	27	80
22	28	88
23	29	96

1 In a preferred implementation, a fast forward trick mode file and a fast reverse
2 trick mode file are produced from an original-quality MPEG-2 coded video main file
3 when the main file is ingested into the video file server. As shown in FIG. 24, a volume
4 generally designated 390 is allocated to store the main file 391. The volume 390 includes
5 an allocated amount of storage that exceeds the real file size of the main file 391 in order
6 to provide additional storage for meta-data 392, the fast forward trick file 393, and the
7 fast reverse trick file 394. The trick files are not directly accessible to clients as files;
8 instead, the clients may access them thorough trick-mode video service functions. With
9 this strategy, the impact on the asset management is a minimum. No modification is
10 needed for delete or rename functions.

11 Because the volume allocation is done once for the main file and its fast forward
12 and fast reverse trick mode files, there is no risk of lack of disk space for production of
13 the trick files. The amount of disk blocks to allocate for these files is computed by the
14 video service using a volume parameter (vsparams) specifying the percentage of size to
15 allocate for trick files. A new encoding type is created in addition to types RAW for
16 direct access and MPEG2 for access to the main file. The new encoding type is called
17 EMPEG2, for extended MPEG2, for reference to the main file plus the trick files. The
18 video service allocates the extra file size only for these files.

For the transfer of these files to archive or to another video file server, it would be useful to transfer all the data even if it is a non-standard format. For the FTP copy-in, a new option is added to specify if the source is in the EMPEG2 format or if it is a standard MPEG2 file. In the first case, the copy-in should provide the complete file 390. In the second case, the video service allocates the extra size and the processing is the same as

a mechanism to achieve a trade-off between the scaled I frame quality and the temporal (motion) sampling. Depending on the speed-up factor (or the target trick mode file size) and also the number of interleaving freeze frames to be inserted, the video service procedure in charge of trick mode file generation determines a sub-sampling pattern (closest to uniform) to choose the original I frames which will be scaled and included in the trick mode files. For example, the case of an original clip with 10 frames per GOP, a trick mode file size which is 10% of the main file together with 0 freeze frames, implies the use of all original I frames for being scaled and included in the trick mode file. This will typically result in a low quality scaling. As another example, the case of an original clip with 10 frames per GOP, a trick mode file size which is 10% of the main file together with 1 freeze frame, implies the use of a 2 to 1 (2:1) sub-sampling on the original I frames which will choose every other original I frame for being scaled and included in the trick mode file.

FIG. 25 is a more detailed diagram of the volume 390, showing additional meta-
data and related data structures. The Inode 401 includes 4 disk blocks containing a file-
system oriented description of the file. The Meta-data (MD) directory 402 includes 4
disk blocks describing each entry of the meta-data area 392. The entries of the meta-data
area 392 include a description of the MPEG-2 meta-data 403, a description of the trick
files header meta-data 404, and a description of the GOP index meta-data 405. The
MPEG-2 meta-data 403 includes 15 disk blocks maximum.

21 The trick files header 404 includes 1 disk block, which specifies the beginning of
22 free area (end of last trick file) in blocks, the number of trick files couple (FF FR), and
23 for each trick file, a speed-up factor, a block address of the GOP index, a block address of

1 Playing a file is done with the CM_MpegPlayStream class. Fast forward
2 (reverse) can only be requested when we are in the paused state. The current frame on
3 which we are paused is known from the MpegPause class. This frame is located in the
4 GOP index of the trick file. Then the clip start point and length are modified in the Clip
5 instance with the trick file position computed from the beginning of the clip. So, the Clip
6 class handle these trick files in a manner similar to the main file. The current logical
7 block number is updated with the block address in the trick file recomputed from the
8 beginning of the main clip. In fact, a seek is performed in the trick file as it was part of
9 the main file, which is totally transparent for the ClipList and Clip classes. The transition
10 from fast forward to pause is handled in a similar fashion. The clip start and length and
11 the logical block number are again updated. The smooth transitions from pause to fast
12 forward and from fast forward to pause are done in the same way as for regular play.
13 There is a splicing from the pause stream to the play stream.

The class hierarchy for trick file handling is shown in FIG. 28. The MpegFast, MpegFastForward and MpegFastReverse class handles the GOP generation from the initial file. This is the common procedure for building the GOP whatever the source and the destination. RealTimeFastFwd and RealTimeFastRev are the class instantiated when a real time fast forward (reverse) has to be done. They manage the real-time buffer flow to the player. There is a derivation of the methods takeBuffer and returnBuffer which uses the base class to build the GOP in the buffer to be played. The main file access is done using a buffer pool.

TrickFilesGenerate is the class instantiated to generate trick files forward and reverse. It inherits from TrickFileAccess the methods for reading the original file some

The bandwidth allocated to the TrickFilesGenerate command is defined in the volume parameters (vsparams or vssiteparams). The selection of a stream server to generate the trick files takes into account this bandwidth only. If preferred stream servers are specified in vsparams (or vssiteparams), then the selected stream server will be one of these specified stream servers.

6 In a preferred implementation of the video service software, a new encoding type
7 is created. The encoding type enum becomes:

```

8      enum encoding_t{
9          ENC_UNKNOWN    = 0,          /* unknown format */
10         ENC_RAW         = 1,          /* uninterpreted data */
11         ENC_MPEG1       = 2,          /* constrained MPEG1 */
12         EMC_MPEG        = 3,          /* generic MPEG */
13         ENC_EMPEG2      = 4,          /* MPEG2 with trick files extension */
14     };

```

```

16 - The encoding information accessible by VCMP_EXTENDEDINFO includes
17 information about trick files:

```

```

19 struct trickFilesInfo_t{
20     ulong_t      generationDate;      /* date/time of the generation of the trick
21     files */
22     rate factor t acceleration;      /* acceleration factor */

```

```

1      ulong_t      framesNumber;          /* frames number in each trick file (FWD and
2          REV) */
3      ulong_t      gopNumber;          /* GOP number of each file */
4  };
5
6  struct EMPEG2info_t{
7      MPEG2info_t      MPEG2info;
8      trickFilesInfo_t      trickFiles<>;
9  };
10
11  union encodingInfo_t switch (encoding-t enc){
12      case ENC_MPEG:
13          MPEG2info_t      MPEG2info;
14      case ENC_EMPEG2:
15          EMPEG2info_t      EMPEG2info;
16      default:
17          void;
18  };

```

20 The video service software includes a new procedure (VCMP_TRICKFILESGEN) for
21 trick file generation, which uses the following structures:

```
23 struct VCMPtrickgenres_t{
```



```

1      rate_factor_t      acceleration;
2      bandwidth_t        reservedBw;
3  };
4
5  cms_status  CMSPROC_GEN_TRICK_FILES (cms_trick_gen_args)      = 34,
6
7  struct trick_gen_completed_args {
8      Handle_t      Vshandle;
9      cms_status      status;
10 };
11 void CTLPROC_TRICKGENCOMPLETED (trick_gen_completed args)      = 8,

```

13 The video service includes the following option to force the regeneration of trick
14 files even if they exist:

```
15 nms content-gentricks <name> [<-f>] [acceleration]
```

Without this option, an error code is returned if the trick files exist. “Acceleration” is an acceleration factor. If it is not present, the default value is taken in vsparams.

The video services include a encoding information access function (nms_content
-m). This function produces a displayed output containing, for each trick file generated,
the acceleration, the generation date and time, the frames number, and the GOP number.

21 For the use of an FTP copy function with the trick files, the following new
22 commands are added:

1 nms_content -copyinfull <same arguments as -copyin>
2 nms_content -copyoutfull <same arguments as -copyout>
3

4 Another application of the SNR scaling of the invention is to reduce the bit rate of
5 an MPEG-2 transport stream in order to allow combining multiple MPEG-2 transport
6 streams to match a target bit rate for a multiple program transport stream. For example,
7 FIG. 29 shows a system for combining an MPEG-2 audio-visual transport stream 411
8 with an MPEG-2 closed-captioning transport stream 412 to produce a multiplexed
9 MPEG-2 transport stream 413. In this case, the closed captioning transport stream 412,
10 containing alphanumeric characters and some control data instead of audio-visual
11 information, has a very low bit rate compared to the audio-visual transport stream 411.
12 Assuming that the target bit rate for the multiplexed transport stream 413 is the same as
13 the bit rate of the audio-visual transport stream 411, there need be only a slight decrease
14 in the bit rate of the audio-visual transport stream, and this slight decrease can be
15 obtained by occasionally removing one non-zero AC DCT coefficient per 8x8 block.
16 Therefore, in the system of FIG. 29, the audio-visual transport stream 411 is processed by
17 a program module 414 for selective elimination of non-zero AC DCT coefficients to
18 slightly reduce the average bit rate of this transport stream. A transport stream
19 multiplexer 415 then combines the modified audio-visual transport stream with the closed
20 captioning transport stream 412 to produce the multiplexed MPEG-2 transport stream
21 413.

22 In order to determine whether or not any non-zero AC DCT coefficient should be
23 eliminated from a next 8x8 block in the audio-visual transport stream 411, a module 421

